

# Chapter 18

## Calorimeter Backend Electronics

### 18.1 Backend Electronics

The calorimeter backend electronics for E989 encompass the systems for the distribution of the clock and synchronization signals to the experiment and for the digitization of the waveforms from each channel of electromagnetic calorimetry. The backend electronics for the tracker and the auxiliary detectors are discussed within those sections. The calorimeter and tracker backend electronics will both be implemented as  $\mu$ TCA Advanced Mezzanine Cards (AMCs).

The backend electronics will also be used to capture the signals from the entrance counters (see Ch. 20), the fiber beam monitors (see Ch. 20), the electrostatic quadrupoles (see Ch. 13), the fast muon kicker (see Ch. 12), and the laser calibration system (see Ch. 17). The channel counts for these systems are

Entrance counter:	2,
Fiber harp:	28,
Quadrupoles:	4,
Kicker:	6-9,
Laser:	42.

#### 18.1.1 Physics Goals

The clock system must provide a frequency-stabilized and blinded clock signal that provides the time basis for the determination of  $\omega_a$  and a second frequency-stabilized clock, tied to the same master clock, for the precision magnetic field measurement.

As discussed in Chapter 16, the calorimeter backend electronics contribute to three fundamental areas in the determination of  $\omega_a$ : the determination of the positron arrival time at the calorimeter, the determination of the positron energy, and the separation of multiple positrons proximate in time (pileup). Given the continuous distribution that the muons reach in the storage ring and the random decay times, waveform digitizers (WFDs) best fulfill these roles. The SiPM output response is deterministic, so a fit to the digitized waveform can determine the arrival time and the energy of each positron and can resolve overlapping signals. The WFDs will digitize each muon fill in its entirety, and the frontend DAQ computers will derive the persistent  $T$  method and  $Q$  method datasets from these continuous

Table 18.1: Summary of the minimum clock and digitization requirements for the calorimeter backend electronics, which are discussed in more detail in Section 18.1.2.

Feature	Driving Consideration	Requirement
Digitization rate	Pileup identification	$\geq 500$ MSPS
Bandwidth	Pileup identification	$\geq 200$ MHz
Bit depth	Energy resolution	12 bits
Station readout rate	Fill length and repetition rate	$\geq 3$ Gbit/s (avg.)
Clock stability over fill	Negligible $\omega_a$ systematic contribution	$< 10$ ps over $700 \mu\text{s}$
Frequency upconversion	Time base for frequency $\omega_a$	$< 1$ Hz
Clock jitter	Signal fidelity for time extraction	$\leq 200$ ps

waveforms. This scheme eliminates hardware-level deadtime and associated efficiencies that would couple to muon intensity and introduce systematic biases that are difficult to control at the sub-part per million level.

The WFD must convert the waveforms from analog to digital while retaining the signal fidelity necessary to meet the calorimetry requirements on energy resolution and pileup differentiation. The system must convert the distributed master clock frequency to the required sampling frequency range while maintaining the timing requirements, without allowing circumvention of the experimental frequency blinding. The digitized waveforms must be transferred without loss to the DAQ frontends for data reduction. The system must also provide the support and infrastructure to capture samples for pedestal determination, gain monitoring and correction, and stability cross checks of the gain monitoring system.

These considerations lead to the minimum requirements summarized in Table 18.1.

## 18.1.2 Requirements

### Clock and synchronization distribution

To avoid systematic biasing of  $\omega_a$ , the distributed master clock must be held stable against systematic phase shifts or timing drifts to under 10 ps over the  $700 \mu\text{s}$  fill [1]. To help maintain signal fidelity and the subsequent extraction of the positron arrival time from fits to the digitized waveform, the random timing jitter should be smaller than the ADC signal sampling window (the ADC's aperture delay), which is of order 100 – 200 ps for the required digitization rates. The frequency upconversion within the WFDs must maintain these requirements, and the upconversion factor must be determined to better than a part per billion.

Synchronization signals such as begin-of-fill must be distributed to each backend electronics channel to signal the time to capture the data. The synchronization signals will flag the specific master clock cycle on which to begin data acquisition for each muon fill. These signals will arrive with the granularity of the master clock. The muon beam entrance-counter signals will be digitized by this WFD system for one of the calorimeter stations and will provide a precise start time measured with the sampling frequency (800 MHz), not master clock (40 MHz) precision. Additionally, the laser system will be pulsed in advance of every fill to dead-reckon the time-alignment of all calorimetry channels.

## Waveform digitization

**Signal requirements:** The energy resolution budget (5% near the 1.5 GeV threshold for fitting) determines the WFD minimum bit depth. Assuming a typical  $3 \times 3$  array of crystals summed to determine the energy, having 8 bits at 1.5 GeV would already contribute 1.2% to the energy resolution. This energy is about half the maximum energy range, and the system should have the overhead for complete study of the pileup energy distribution, which requires 10 effective bits. The effective number of bits is typically between 1 and 2 bits lower than the physical ADC bits. A digitization depth of 12 bits provides the appropriate resolution.

The signal separation characteristics will be determined by a combination of the crystal wrapping<sup>1</sup>, SiPM and amplifier response (see Figure 17.16), total cable, and WFD bandwidth. The WFD bandwidth must be large enough to avoid significant stretching of the pulse shapes, with the rise time remaining under 2 ns (if the final wrapping choice allows). The overall pileup requirement for the experiment (see Section 17.1: the system must be able to distinguish pulses separated by 5 ns) drives the specification for the digitization sampling rate. Laboratory tests at 800 megasamples per second (MSPS) show clear separation of pulses with this separation (see Figure 17.17). Figure 18.1 shows the fits for two pulses measured in the lab with 5 and 10 ns separations at 500 MSPS sampling rate. We can clearly resolve even the 5 ns separation at 500 MSPS, but we will lose fidelity at lower sampling rates. A higher digitization rate will give more headroom to cope with higher intensity muon beams, and there are 800 MSPS ADCs available that afford a good balance between performance and cost. Furthermore, with the improved SiPMs now in hand, their inherent signal characteristics are much faster, with rise times of order 2 ns and under 4 ns fall times. These times also push us towards the higher sampling rate. Simulations based on the Čerenkov photons show that, at 800 MSPS, the presence of a second soft cluster can be identified 100% of the time down to a time separation of 3.5 ns. With the additional factor of 3.0 that comes from the  $\text{PbF}_2$  segmentation, this performance can accommodate beam intensities of a factor of 2.5 over the proposed intensity (10.0 over the Brookhaven E821 intensity). Our baseline implementation will be based on the 800 MSPS rate.

**Physical requirements:** The WFD crates will be located just over 1 m from the dipole field of the storage ring, where the fringe field ranges from 30 to 60 Gauss. Ideally, perturbation of the storage ring field by the resulting magnetization of the materials in each WFD station would be limited to well under a part per million. We can roughly limit the acceptable level of a magnetic material by assuming that a magnetized sphere of that material in a uniform fringe field should at most create a static perturbation of 0.1 part per million. A predominantly Aluminum chassis will be no problem: 15 kg would result in a perturbation under 0.1% of this limit. For ferromagnetic materials, however, the total mass must be kept under about 200 g, which may require the power supplies to be located farther away. An Aluminium  $\mu$ TCA crate was shipped to Fermilab during summer 2014 for characterization using a test magnet. The results of this test revealed the need to change two pieces of the crate that have been addressed to VadaTech. The first piece is the metal connector between the crate power output and the power module that delivers power to the modules in the crate. This connector will from now on be made out of plastic. The second

---

<sup>1</sup>The GEANT4 Cherenkov light simulation (Section 17.4.1) shows that the wrapping material creates rise time contributions that vary from 1 to 3 ns depending on the material used.

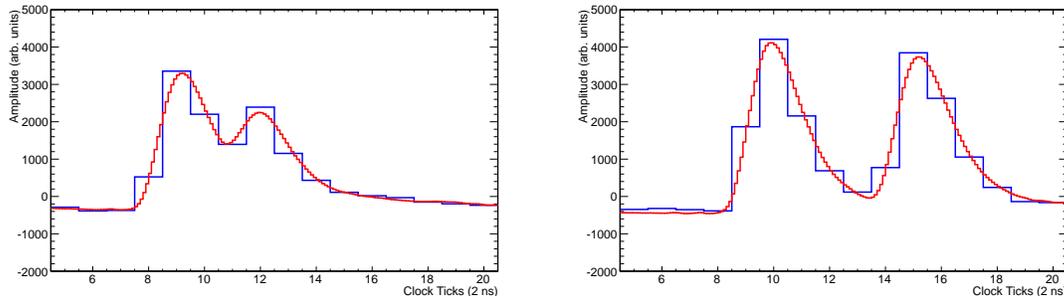


Figure 18.1: Fits to test data (see Section 17.4.1) with two pulses separated by 5 ns (left) and 10 ns (right) with a 500 MSPS sampling rate. The fits (red) clearly resolve the two peaks in the data (blue), even for the 5 ns separation.

piece is the steel card cage within the  $\mu$ TCA crate that will be made of Aluminium. Final assessment will be made after the final version of the  $\mu$ TCA crate from VadaTech arrives at Fermilab during summer 2015.

**DAQ requirements:** During experimental running, muons will be stored in the storage ring for 700  $\mu$ s fills. The basic fill structure will be two groups of eight fills, with the fills within a group occurring at 10 ms intervals and with the groups of eight occurring at 197 ms intervals (see Figure 7.2). This basic structure repeats every 1.4 s for an average fill rate of 12 Hz.

To eliminate dead-time, up to a 700  $\mu$ s waveform for each calorimeter channel will be digitized and transferred to the DAQ frontend system for data reduction. Each WFD station must provide adequate buffering and throughput to support the average data rate. Assuming an 800 MSPS digitization rate, the average rate is 5.9 Gbit/s for 12-bit samples being transmitted as 16-bit words and 4.4 Gbit/s if bit-packed, i.e., 12-bit samples being transmitted as 12-bit words. The experiment may run with a shorter 600  $\mu$ s sampling time per fill, which would reduce the data throughput to 5.1 Gbit/s (not bit-packed) and 3.8 Gbit/s (bit-packed) for each station.

### 18.1.3 Baseline Design

#### Clock distribution

The clock system will distribute a high-precision clock and synchronization signals to each backend DAQ crate. It will provide an external time reference that is fully independent of accelerator timing to ensure that the  $\omega_a$  measurement is not biased by synchronous accelerator or ring events.

The clock system will utilize a series of off-the-shelf components in conjunction with a fiber optic distribution and encoding system that was developed for the CMS experiment [5]. An optical clock distribution system minimizes concerns due to pickup and ground loops. Overall design and consideration of the clock system follows from experience gained in Brookhaven E821 [2] and MuLan [3].

A block diagram of the clock system is shown in Figure 18.2. The principal clock source for both the  $\omega_a$  and  $\omega_p$  measurement will be produced by a Meridian Precision GPS TimeBase, a GPS-disciplined oscillator. The utilization of the clock in the  $\omega_p$  measurement is described in Section 15.2.4. The Meridian module is supplemented by a “low-phase-noise” output module to minimize jitter. The GPS system additionally provides time-stamps. The GPS clock produces a 10 MHz output signal which is fed to a Stanford Research Systems SG380 series RF signal generator providing a shifted  $\omega_a$  clock of 40 MHz plus a small offset ( $\epsilon$ ) that will be blinded.

The  $40 + \epsilon$  MHz clock is then carried on double-shielded RG-142 cable from the signal generator to the encoding and fanout system, which resides in a single  $\mu$ TCA “clock crate.” The analog clock signal is injected into an FC7 board [5] equipped with a standard FCM DIO 5CH TTL A module [6]. The FMC card will additionally receive input signals from the Fermilab accelerator and other Muon  $g-2$  subsystems. To define the “begin of fill” signal, we have identified the “recycler beam sync extraction event to muon” signal. This signal will be distributed to all backend crates and, after time-alignment, will synchronize and initiate data acquisition (i.e., provide a common start.) The other control signals received by the FC7 will be utilize to define run mode and synchronize the DAQ across systems.

With the 40 MHz clock and control signals in the FC7, the clock frequency is multiplied by four and the control signals are encoded via the TTC protocol developed for the Large Hadron Collider (LHC) experiments [8]. The encoded clock and controls will then be sent via optical fiber to an AMC13 card working in dual-star mode within the clock crate. The AMC13 will distribute the recovered clock and controls along the  $\mu$ TCA backplane to two fanout FC7 boards which, in turn, will each produce 16 copies of the encoded clock and control signals. These outputs will carry signals to the backend  $\mu$ TCA crates.

A single fiber carrying the clock and control signals is distributed to the AMC13 board on each backend  $\mu$ TCA crate. The AMC13 receives the signal, recovers the base 40 MHz clock, and decodes the control signals onto a parallel stream. The clock and control signals are then sent along the backplane to each of the WFDs. On the WFD, the clock frequency is upconverted to the 800 MHz base sampling rate.

The clock crate will drive a single clock/control fiber to each of the 24 calorimeter backend crates, 1 tracker crate, 1 auxiliary detector, 1 laser calibration system crate, and the DAQ system.

Monitoring of the clock system will occur at several stages. A second Rubidium oscillator will reside in the same crate to check the master clock stability. The stability of the difference between master and reference clock frequencies plus a blinded offset will be monitored directly in the counting room, using a data login featured oscilloscope. At the receiving end, the AMC13 will verify clock functionality with an internal counter compared to a local oscillator. Further, direct tests on time slewing and other systematic effects will be performed using the clock signals as seen by the WFDs.

To date, the individual components (GPS-disciplined clock source, frequency synthesizer, AMC13, backplane distribution) have all been tested. In the near future, we will assemble a complete system in order to begin integration testing.

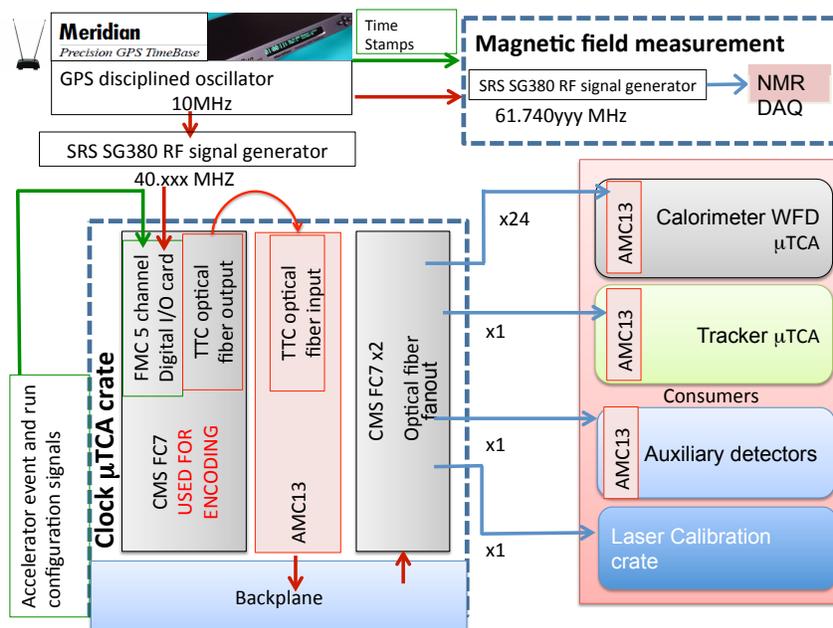


Figure 18.2: Block diagram of the clock distribution system.

### Clock and Commands Center (CCC)

As described above, three FC7 boards will be utilized in the clock crate: one will encode the clock and command signals using the TTC protocol (the Clock and Commands Center, CCC) and two will be used to fan the TTC signals out to backend crates via optical fiber. The CCC FC7 board will receive the clock and relevant timing signals from other systems. Logic on the CCC will then translate the signals into a valid TTC transmission, where the proper command bits are encoded along with the distributed clock. For the specific example of a normal run mode, where muons will be injected into the Muon  $g-2$  storage ring, Figure 18.3 shows the approximate timing for the command logic.

The CCC will receive the begin-of-fill (BOF) signal which will be utilized to initiate data acquisition operations. The BOF signal will arrive at the CCC about  $3 \mu\text{s}$  prior to beam entering the storage ring. This time is sufficient for logic, handshaking, and distribution of the relevant signals. Within  $1 \mu\text{s}$  after the BOF is received, the CCC will acquire the “DAQ ready” signal and Laser Calibration System inputs. To initiate data acquisition in the backend crates, we will require both BOF along with the DAQ-ready signal. If the DAQ is not ready to acquire data, the backend electronics will not acquire data. Since our DAQ system is designed to keep pace with data acquisition at all times, a DAQ not-ready state at begin-of-fill time is likely to indicate a problem that needs to be address. The logic “AND” between DAQ ready and begin-of-fill signal will set the trigger of the experiment (A CHANNEL bit of TTC signal will be set to 1). This operation will take about 100 ns after all the inputs have been received and checked. Additional information can be encoded and

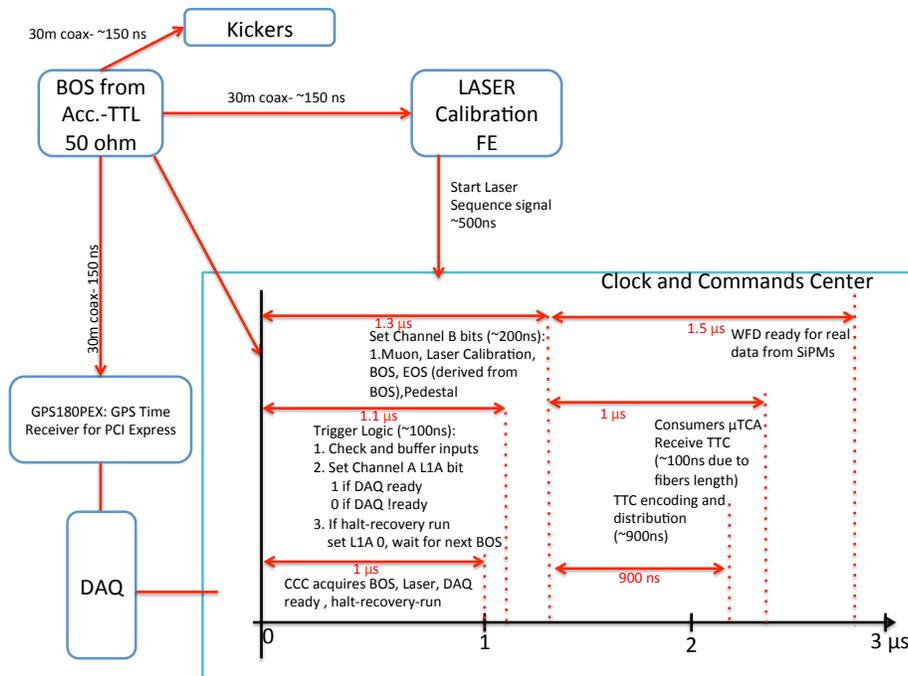


Figure 18.3: A rough estimate of the trigger and commands center time scale

sent to the backend through the B CHANNEL TTC signal. At this stage the A CHANNEL and B CHANNEL info will be encoded and distributed. The final TTC signal will be then received by all of the consumers after about  $2 \mu\text{s}$  from the real BOF arrival time. This signal will be delayed and time-aligned as needed so that all backend consumers receive the information at a well-defined time prior to the arrival of beam.

The system outlined here provides a “common start” to the DAQ system, but is not the source of the precision arrival time of the beam. Precise begin-of-fill timing will be provided by the entrance counter signal.

Examples of other run types include laser and pedestal runs. These run types could be initiated synchronously with beam or between muon fills. The CCC system will also be able to initiate a global reset signal to all backend crates which can be utilize to resynchronize the data acquisition process.

### Waveform digitization

The proposed system draws heavily on the hadronic calorimeter and DAQ upgrade [4] underway for the CMS experiment at the LHC, which uses  $\mu\text{TCA}$  technology. The WFDs for each calorimeter station will reside in a single Aluminum VadaTech VT892  $\mu\text{TCA}$  crate as a set of 12 five-channel AMCs. The first five Aluminum VT892 crates have been delivered in May 2014. The first 10 modified  $\mu\text{TCA}$  crates, which resulted from the magnetic field testing at Fermilab, should be delivered by summer 2015, and the currently owned crates will be sent back for upgrade. The remainder of the crates will be acquired during the project implementation phase.

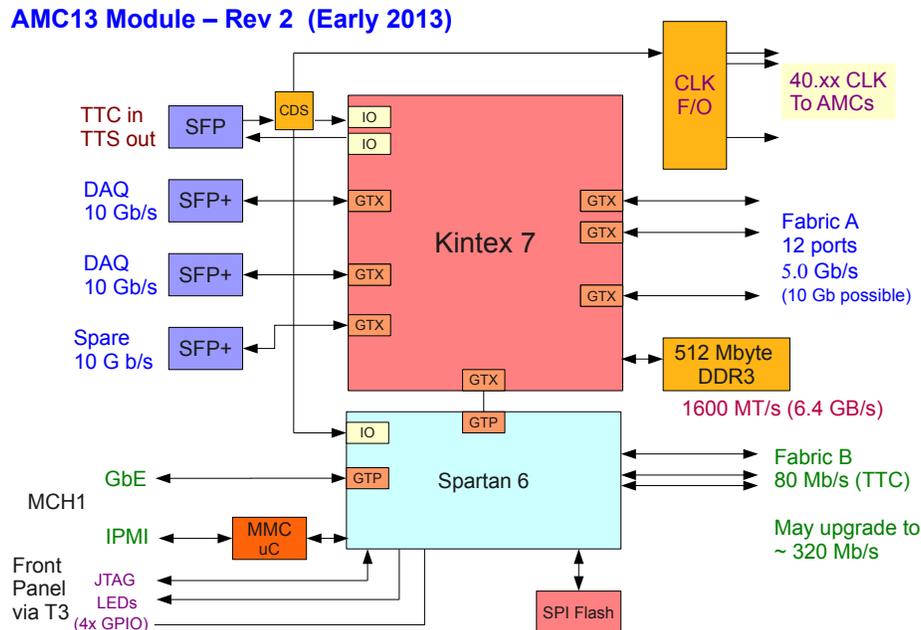


Figure 18.4: Block diagram of the CMS-designed AMC13  $\mu$ TCA card that will control the  $g-2$  WFD readout.

A VT892 crate accommodates 12 full-height AMC cards. Eleven AMCs will instrument the 54 calorimeter channels for one calorimeter station. The 55<sup>th</sup> channel will be used for the laser monitoring source. The 12<sup>th</sup> WFD AMC will reside in the crate, which leaves five calorimeter hot spares per station. The  $\mu$ TCA choice brings a robust system designed for remote operation and monitoring with cooling, power distribution, and clock distribution capabilities already engineered. There will be a separate VT892 crate for each of the 24 calorimeter stations.

**$\mu$ TCA infrastructure** The heart of the CMS  $\mu$ TCA data acquisition system is the Boston U. (BU)-designed AMC card, the AMC13 [7], which replaces a second (redundant)  $\mu$ TCA Carrier Hub (MCH) in the  $\mu$ TCA crate. The AMC13 is shown in block diagram form in Figure 18.4. The AMC13 has the responsibility for distribution of external clock and synchronous control signals to the AMC cards, for readout and aggregation of the data for each fill from the AMC cards, and for transfer of the data to the DAQ frontend computers over one to three 10 Gbit/s optical ethernet links. The AMC13 has gone through an extensive prototyping and testing cycle in CMS, and BU has recently overseen the first large production run of AMC13 modules. CMS has demonstrated that the AMC13 is compatible with the MCH from both VadaTech and N.A.T. Based on feedback from CMS related to reliability and support, we will deploy the VadaTech MCH in our  $\mu$ TCA crates. Figure 18.5 shows the Cornell U. test stand populated with the AMC13, a VadaTech MCH, and one of our Revision 0 five-channel WFD prototypes.

The  $\mu$ TCA backplane connects each of the 12 WFD AMC cards to the AMC13 and MCH

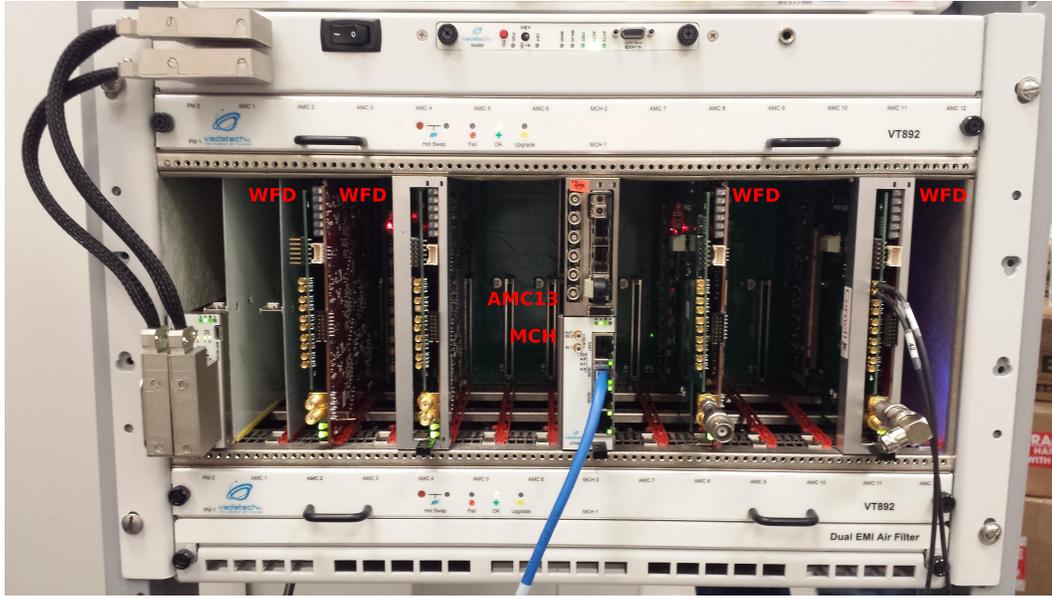


Figure 18.5: The Cornell U.  $\mu$ TCA test stand based on the VadaTech VT892, populated with a CMS AMC13 (upper of center pair of modules), a VadaTech UTC002 MCH (lower of the central pair), and one of the Revision 0 five-channel WFD prototypes (at right) under testing.

in a star topology (Figure 18.6). High-speed readout occurs over Fabric A and is managed by the AMC13's Kintex-7 FPGA. The star topology and AMC13 implementation allow for parallel readout of the 12 WFD cards at rates up to 5 Gbit/s per link (60 Gbit/s aggregate). For continuous readout at the average 12 Hz muon fill rate, only a 0.6 Gbit/s link to each card is necessary for an 800 MSPS sampling rate with the 12-bit samples being transmitted as 16-bit words. There is plenty of overhead in the backplane transmission rate to support the average readout rate. The system can also accommodate readout in the smallest inter-fill time (10 ms), which requires a link rate of 4.5 Gbit/s assuming transfer of 16-bit words, or 3.4 Gbit/s if bit-packed. The WFD system therefore has the flexibility to support a wide range of beam delivery options.

The AMC13 receives the clock and synchronous signals from the clock system via optical fiber using the TTC protocol [8]. With the 40 MHz master frequency, the clock and synchronous TTC data are distributed as a 160 MHz biphas-encoded signal. The AMC13 recovers the data and clock from the optical signal, routes the TTC data over Fabric B via the Spartan-6 FPGA, and fans out the clock over the CLK1 fabric (see Figure 18.7.) On each WFD AMC card, a Texas Instruments (TI) clock synthesizer chip (LMK04906) will upconvert the master frequency to the 800 MHz frequency needed for digitization.

The AMC13 event builder has been upgraded from its specific CMS HCAL implementation to support both Muon  $g-2$  and broader CMS deployment [9]. In particular, the event builder supports much larger event block sizes than it did in its original 4 kB HCAL implementation. A beta version of this new firmware, including the communications firmware block for the AMC side of the 5 Gbit/s link, has been delivered to Cornell U. in spring 2014

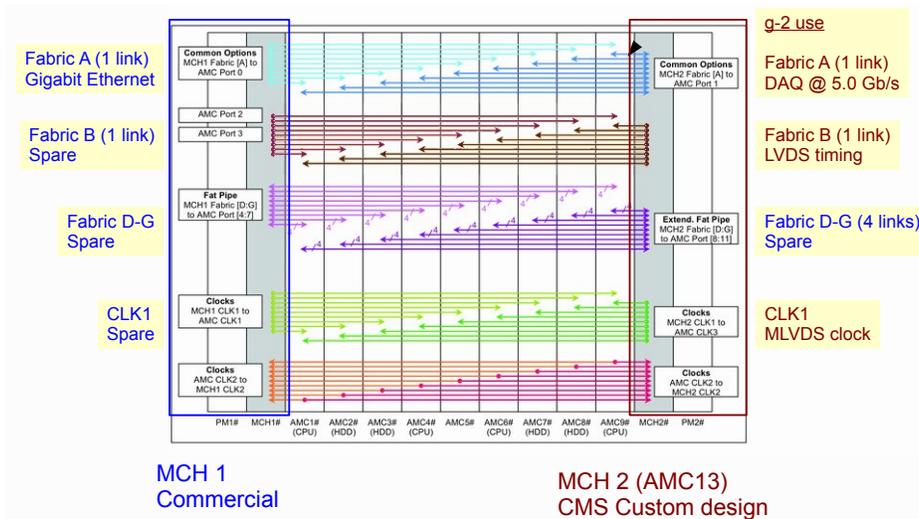


Figure 18.6: Dual-star backplane configuration for use with  $g - 2$  AMC13. High speed data transfers proceed over Fabric A. Timing and synchronization proceed via Fabric B.

and has been shown to work. Several upgraded versions have been successfully tested since then.

Each AMC13 communicates the event data to its corresponding DAQ frontend computer via one to three 10 Gbit/s TCP/IP optical links. A Muon  $g-2$  specific modification of the AMC13 firmware that supports TCP/IP has been released to Cornell U. and U. of Kentucky for testing. The implementation will use the 512 MB memory buffer on board the AMC13, which can accommodate eight fills of data without compression. Significantly more buffering is available on board the WFD AMCs themselves, as discussed below. With the buffering, the AMC13 can communicate asynchronously with the external DAQ system to support the average data rate of about 6 Gbit/s per station. Only one of the three 10 Gbit/s optical links for communicating with the DAQ frontend computers is therefore needed.

The AMC13 has substantial on-board processing capability in the Kintex-7 FPGA that could be used for a variety of tasks. These tasks could include lossless encoding/decoding of the data on-the-fly to allow deeper buffering of the data or to prepare the  $T$  or  $Q$  method datasets. The latter would substantially reduce the required bandwidth between the AMC13 and the DAQ frontends as well as relieving processing requirements on the DAQ, if needed.

**Waveform digitizer hardware** The baseline WFD design is centered on the TI ADS5401, an 800 MSPS 12-bit ADC with an input bandwidth of over 1.2 GHz. The block diagram for the anticipated final version of the five-channel AMC card is shown in Figure 18.8 and the Revision 0 of the five-channel prototype based on the 800 MSPS ADS5401 is shown in Figure 18.9. The WFD AMC has three boards. The main board is the heart of the system, with the ADCs, buffer memory, channel and fabric FPGAs, clock synthesizer, and  $\mu$ TCA management controller (MMC). The power supply and regulation, which takes the 12 V payload power distributed by the crate and converts it to the voltages needed on board, is implemented on one of the mezzanine cards. On the one-channel prototype (based on the 500 MSPS ADS5463 from the conceptual design), this block was at the lower right on the

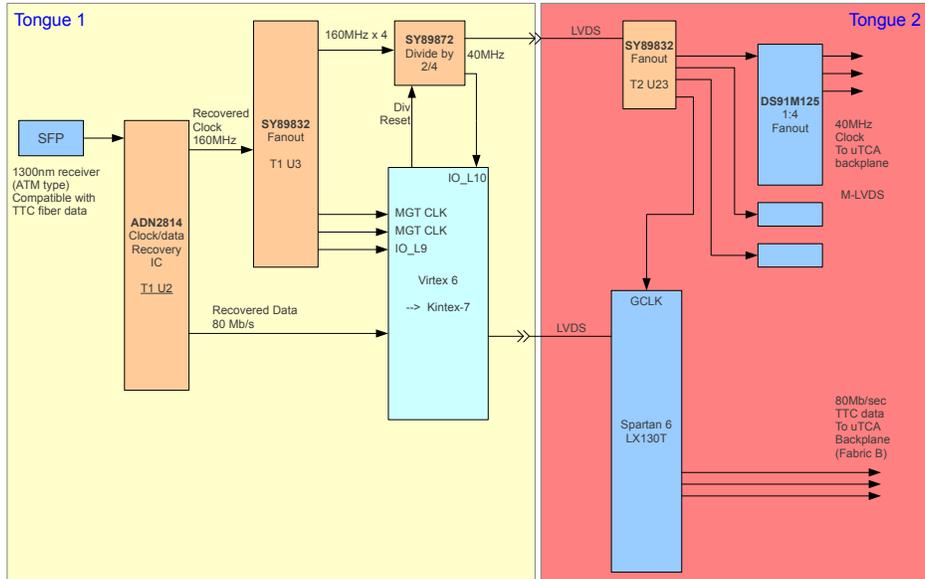


Figure 18.7: Timing paths through the AMC13 for the LHC. The  $g-2$  experiment will utilize the FPGA-free LVDS path, which distributes the 40 MHz master clock.

main board, but it suffered from power regulator failures that rendered the board irreparable. In addition to switching to a different regulator, we have moved the power supply to a daughter card. If a regulator fails, we can easily replace the daughter card. In addition, this design freed up space on the main board to accommodate all five WFD channels. The third board contains the analog frontend (AFE) and connects to the main board via shielded connectors. This design gives us the ability to fine tune the AFE design without affecting the main board where the cost driver components reside. The design also provides flexibility for future deployment. In the baseline design, the AFE has a differential input impedance of  $100\ \Omega$ , with a maximum 2 V peak-to-peak differential signal, and will bandwidth limit the signal to 200 MHz.

On the main board, each channel has a dedicated Kintex-7 FPGA (XC7K70T-2FBG484C) that controls the data flow out of the ADC to a dedicated  $64\text{M} \times 16$  bits SDRAM memory buffer as well as the streaming of data over a high-speed serial link to the fabric FPGA. At an average 12 Hz fill rate, the SDRAM can buffer over nine seconds of data with no bit-packing. The fabric FPGA is another Kintex-7 (XC7K160T-1FBG676C), which controls the data streaming out of each channel and over the backplane Fabric A to the AMC13 event builder. Both the channel to fabric FPGA link and the fabric FPGA to AMC13 link operate at 5 Gbit/s. The data will transfer out of the five channels sequentially, so the data from all cards can be transferred to the AMC13 in just under 9 ms for the 800 MSPS 12-bit baseline. The fabric FPGA also supports the 10 Gbit/s ethernet link and IPbus [10], an IP-based protocol for controlling  $\mu\text{TCA}$  hardware devices.

The AMC card must frequency lock on the blinded master  $40 + \varepsilon$  MHz clock and upconvert it to an  $800 + \alpha \cdot \varepsilon$  MHz clock for the ADCs, where  $\alpha$  is the upconversion factor between the master and ADC sampling clock. The WFD cards will receive the 40 MHz clock via the

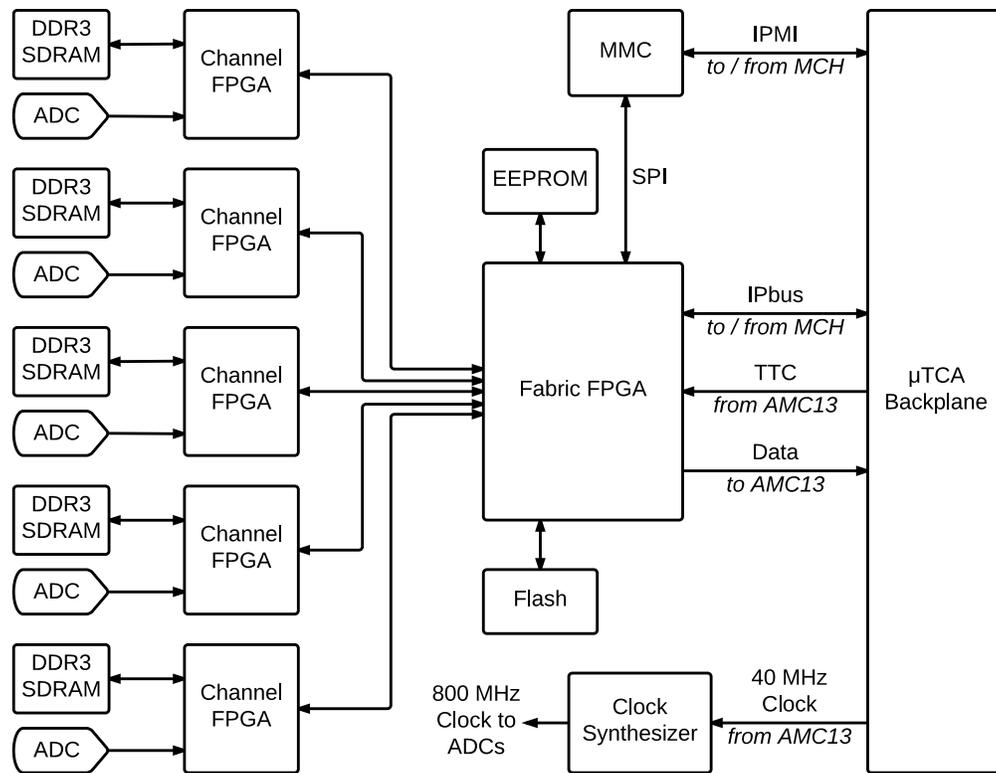


Figure 18.8: Block diagram of the anticipated final version of the five-channel WFD AMC card.

$\mu$ TCA backplane. The clock is distributed by the AMC13 via the FPGA-free LVDS clock path shown in Figure 18.7. The full clock path will need significant testing to verify that it will have a highly stable duty cycle, slew and wander within the phase stability specifications over a fill, and no differential nonlinearities. We have, however, consulted closely with the engineer responsible for the clock for operation of the Cornell Electron Storage Ring (CESR), which also has stringent timing requirements. He does not anticipate any intrinsic difficulty in meeting the  $g-2$  stability specification. On timescales of several hundred  $\mu$ sec, a more important issue is typically environmental noise. We must ensure, for example, that the clock supplied to the ADC will be immune to noise sources, especially those correlated with the fill structure such as the firing of the kicker. Because environmental noise is an issue, we use a single package clock management, based on the TI LMK04906, rather than a discrete component solution. This chip can upconvert the 40 MHz input clock to the 800 MHz range and distribute it over five output channels with a programmable delay on each line. The programmable delay will allow correction of channel by channel timing differences in signal path lengths from the photodetectors at the sub-clock-cycle level. The single package device will be far more immune to external noise than a discrete component solution and has much better overall jitter specifications (under 200 fs) than a discrete component solution. The detector-related systematic error budget of 70 parts per billion for the muon spin precession sets the error budget for the ADC sampling clock frequency to be of order 1 part per billion, i.e., 1 Hz. Extensive tests of the upconversion capability and stability of the TI LMK04906

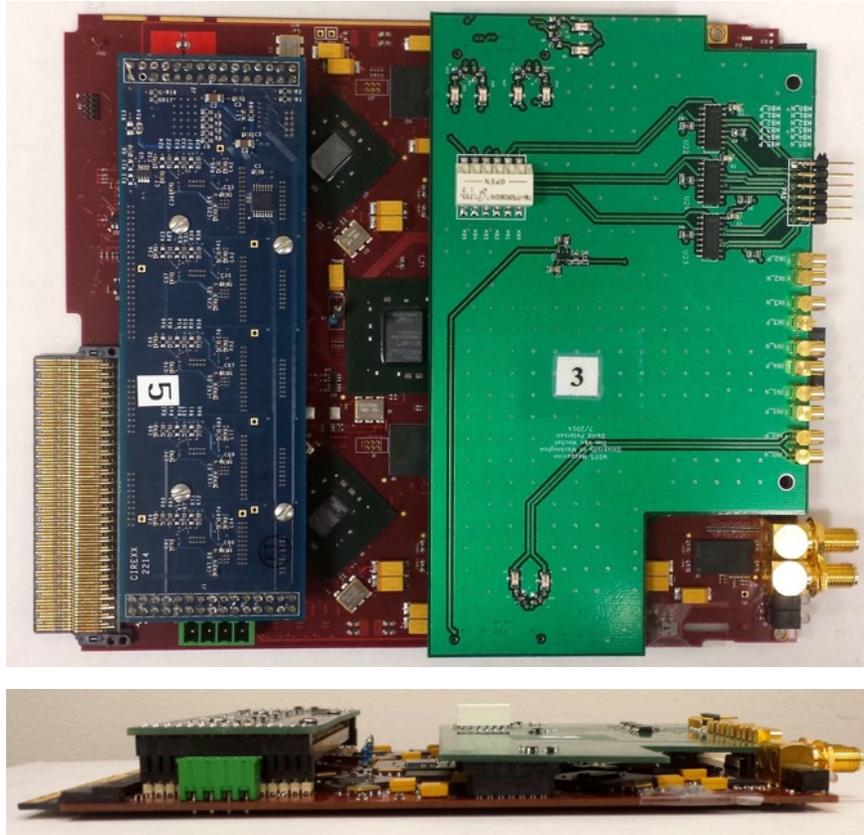


Figure 18.9: The five-channel Revision 0 prototype of the WFD AMC card: the main board is in red, the power supply in blue, and the analog frontend in green.

were performed at Cornell U. and showed that it meets the requirement.

The full clock chain is still undergoing testing to ensure the entire path, from the GPS-stabilized master clock through to the individual ADC channel, satisfies the clock requirements. If not, two backup paths for clock distribution exist that can be used, depending upon the weak link. In one option, the clock could be input to a CERN-developed version of the AMC13 Tongue 3 over copper, independently of the optical TTC signal. This path would bypass the TTC encoding that creates the optical signal as well as the local clock/data recovery on the AMC13. The AMC13 Tongue 2 would still distribute the clock via M-LVDS over the  $\mu$ TCA backplane. In the other option, the master clock would be fanned out externally and input directly into the front panel of each WFD AMC.

To take full advantage of CMS  $\mu$ TCA development, we use the same Atmel microcontroller as CMS to implement the AMC MMC. This same microcontroller has been deployed on the CMS HCAL AMC cards as well as on the AMC13 itself.

At the time of writing, the Revision 0 of the five-channel prototype (main board and power board) that succeeded the one-channel prototype has been extensively tested and has met the baseline design goals. The last prototype, the Revision 1 WFD, was designed taking advantage of the Revision 0 testing to implement new capabilities and fix non-critical issues. This Revision 1 is currently under production and assembly and will be tested summer 2015.

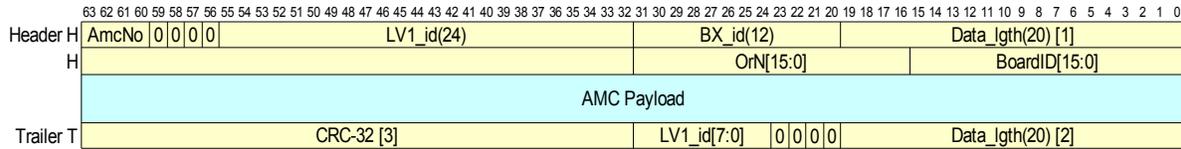


Figure 18.10: The data format for delivering the WFD AMC event payload to the AMC13 event builder.

We will incorporate any modifications (expected to be minor) from the Revision 1 board into the final design, which will go to full production in fall 2015. An AFE board has been provided by U. of Washington for the July 2014 SLAC test beam in order to allow waveform to be sent to the five-channel WFD. This AFE was used to test the WFD Revision 0 at Cornell U. Recently, a more elaborate AFE board was designed at Cornell U. to replace the previous one. This new AFE takes a DC-coupled input signal from the SiPM frontend board via ECDP ribbon cables and offsets this signal using a digital-to-analog converter driven by the main board fabric FPGA. This digital offset will allow utilization of full input range of the ADC. This version was sent for production and assembly.

**Waveform digitizer firmware** The overall operational scheme for the WFD AMC is conceptually straightforward, and makes direct use of the development efforts for the CMS  $\mu$ TCA environment. IPbus is used to configure the WFD AMC operational mode and its associated registers, and it operates over the gigabit ethernet link managed on the Fabric A connection to the MCH (see Figure 18.6). The TTC protocol provides synchronous configuration and triggering over Fabric B. For example, begin-of-fill signals<sup>2</sup> are delivered synchronously to the WFD AMCs to trigger data collection. Through the TTC broadcast command facility, the WFD can be configured to acquire data for a standard muon fill, a pedestal calibration, or a laser calibration event, which differ in their total sampling times. The sampling times for each data type will be configurable via a corresponding FPGA register.

When the AMC13 receives a begin-of-fill signal over its optical input, it relays it over Fabric B to the AMCs and buffers the begin-of-fill in a FIFO. After building a given event, the AMC13 then reads the next buffered fill from each AMC or waits for that data to become available if it is not yet captured. When a WFD receives the read request, the fabric FPGA will control streaming of the data sequentially from each of the five channel DDR3 memory buffers via the corresponding channel FPGA. The fabric FPGA then packages its event data in the data format expected by the AMC13 event builder, which is shown in Figure 18.10.

Figure 18.11 shows the block diagrams for the fabric and channel firmware designs. The main functionality is now largely implemented and under testing on the five-channel prototype WFD AMCs. The firmware development again takes full advantage of available industry-standard designs and other high-energy and nuclear physics (CMS, in particular) development efforts. BU provides the firmware block for the backplane link to the AMC13. The IPbus firmware is available from its open hardware project. We have implemented our high speed channel FPGA  $\leftrightarrow$  fabric FPGA serial link using the Xilinx Aurora link-layer pro-

<sup>2</sup>The  $g-2$  begin-of-fill signals are the equivalent of the CMS “level-1 accept” signals.

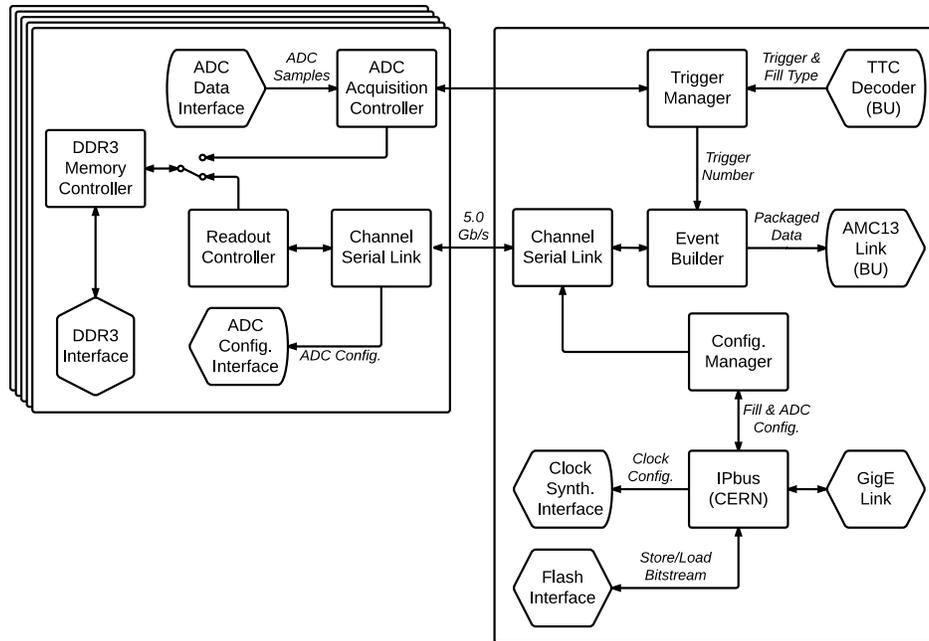


Figure 18.11: Firmware block diagrams for the WFD fabric FPGA (right) and the individual channel FPGA (left) at the time of the Revision 0 five-channel WFD.

to col. The remaining control architecture is built on Xilinx’s implementation of the AXI4 protocol.

The fabric and channel FPGAs have the ability to load their firmware from onboard flash memory at power up. Two “personalities” for the WFD operation will be available for loading. The first one is the standard  $g-2$  mode that includes data taking of a 700  $\mu\text{s}$  fill, laser pulses for calibration and pedestal noise. A TTC command will be issued that causes the WFD to change its acquisition parameters for the three aforementioned data taking type. The second personality is a self- or externally-triggered mode for cosmic ray running. The flash memory is remotely reconfigurable to allow straightforward updating of the WFD installation for bug-fixes, alternate personalities, etc.

**Auxiliary detectors and laser calibration system** Since the previous TDR version in July 2014, the decision was made to use the WFD AMC designed for the calorimeters to capture signals from all the auxiliary detectors (entrance counters, fiber beam monitors, fast muon kicker, quadrupoles) as well as from the laser calibration system. The WFDs for the auxiliary detectors will be gathered in the dedicated “auxiliary detector”  $\mu\text{TCA}$  crate, with the exception of the two WFDs for the entrance counters located in one calorimeter crate. Both the fiber harps and the fast muon kicker should use about a 200 MSPS digitization rate, and the quadrupoles should use about a 20 MSPS digitization rate. The TI LMK04906 clock synthesizer can accommodate those rates, and it was shown at Cornell U. that the WFD can work reliably at those rates. The entrance counter signal will be digitized at 800 MSPS, and the two corresponding WFD channels will be located in one of the calorimeter  $\mu\text{TCA}$  crates. The auxiliary detector  $\mu\text{TCA}$  crate will contain 12 WFDs, i.e., 30% of them being hot spares.

The laser calibration system (see Section 17.4.3) will require one hot spare channel of WFD per  $\mu$ TCA crate (the 55<sup>th</sup> channel) for monitoring the laser pulses at the detector end. A separate  $\mu$ TCA crate in the laser hut will host additional WFDs for monitoring the source laser pulses. The nominal WFD digitization rate of 800 MSPS will be used.

### 18.1.4 Performance

The overall system is designed to operate to stream at the 5 Gbit/s line rate of the FPGA GTX serial lines. Under nominal running conditions, the required average data transfer rate for a five channel card is just under 0.6 Gbit/s, so the design has almost an order of magnitude of headroom. Cornell U. has established the WFD AMC  $\leftrightarrow$  AMC13 backplane link and recorded digitized data coming out of the AMC13, sending an analog waveform to the Revision 0 five-channel WFD AMC prototype. Extensive tests performed recently meet the design requirements. The final WFD prototype was sent for production and assembly and will be tested over summer 2015, leading to the final production of 308 WFD AMCs (15% of hot spare channels).

Simulation and test beam measurements indicate that the proposed baseline meets the basic energy, pulse separation, and random jitter requirements. As noted above, the level of control of systematic timing trends over the 700  $\mu$ s fill times with the expected fill structure must still be characterized, but we do not expect a serious problem.

The proposed solution also provides the experiment with significant flexibility. Should the opportunity arise, for example, for a higher average rate of muon fills, there is no intrinsic limitation from the  $\mu$ TCA-based solution outlined here.

Based on measurements of the power draw of the WFD and the other modules in the crate, we expect each station to consume approximately 420 W of power, which is safely below the maximum power of the 962 W power module available for the VadaTech crate.

The  $\mu$ TCA solution also provides the platform for the tracker readout boards (see Section 19).

### 18.1.5 Value Engineering and Alternatives

The choice of the  $\mu$ TCA platform and the CMS AMC13 continue to represent significant value engineering. The  $\mu$ TCA platform comes with timing, cooling, power, mechanical, and remote monitoring elements all pre-engineered into the system. We found the engineering cost of a custom solution, such as one based around less expensive PCIe technology, would result in a similar total cost but with higher associated risk. The AMC13 choice for data readout with the  $\mu$ TCA solution has revealed several examples of value engineering. For example, we originally anticipated that  $g-2$  would need custom AMC13 event builder firmware. However, with the broader dissemination of the AMC13 throughout the CMS sub detectors, CMS required similar modifications to the original AMC13 event builder firmware. The BU group engineered a common solution that works well for both experiments, halving the development cost for each experiment. Similarly, we have been able to adopt the IPbus control firmware and the Atmel microcontroller MMC firmware, and both have functioned well “out of the box” with only minor modifications for the  $g-2$  environment.

We also considered COTS waveform digitizers. When approaching Struck, though, the reply we received was “For an application in the 1500 channel count I tend to assume, that a custom card may be advised to optimize performance and cost to the application.” We continued on the path of developing our own.

Given both the increased signal speed in the SiPM response since the CDR and the increased anticipated muon flux, we have moved to a baseline design centered on the 12-bit 800 MSPS ADS5401 ADC from Texas Instruments. This design required a faster speed grade FPGA than the original design based on the TI ADS5463, a 12-bit 500 MSPS.

We have investigated discrete component clock circuitry on the WFD AMC card based on the AD9510 or another similar clock synthesizer. The design included a clock delay line for each channel that can correct for differences in signal path lengths from the photodetectors at the sub-clock-cycle level. Such a design would have significantly more inherent jitter (several tens of picoseconds) and would have greater sensitivity to environmental noise.

For the clock system, we have considered alternate clock sources. The clock for Brookhaven E821 was disciplined by the LORAN-C signal which is now obsolete. Undisciplined Rubidium oscillators would likely deliver the precision necessary; however, the GPS disciplined oscillator provides long term stability as well as additional features like time-stamps, which are of particular use to the field measurement. Options for the frequency synthesizer are limited by the 40 MHz signal necessary for the WFD.

### 18.1.6 ES&H

The  $\mu$ TCA crate for each calorimeter station will weigh approximately 30 pounds and will be supported by the calorimeter housing (see Section 17.6). Power to each crate will be supplied by a 60 – 70 V supply that connects to an in-crate power module that maintains the stable 48 V on the backplane. When fully populated with the WFDs, the each station will draw approximately 500 W of power. If the magnetic field requirements allow, the power supply will be resident on the crate. If not, the supplies will be located more centrally in the ring, with a few meter cable run and the supply voltage closer to 70 V.

The latter configuration in particular involves high voltage with several amps of current. We will ensure that all our equipment and installation conforms to the Operational Readiness Clearance criteria.

### 18.1.7 Risks

The largest risk in the WFD project arises from the distribution of the optical clock signal through the AMC13 and  $\mu$ TCA backplane and, in particular, whether that path will meet the frequency and phase drift requirements. To mitigate risk to the project, the WFD AMC cards are designed to allow timing and synchronization inputs via the front panel. This allows for a “brute force” technique virtually identical to that employed in Brookhaven E821 and MuLan, where multiple fanouts were used to distribute analog clock signals directly to each WFD front panel. This alternative would also require a modified clock fan-out and cabling scheme. The total differential cost to the experiment should be under \$40k for engineering and production. Biases in the clock translate directly into biasing of  $\omega_a$ , so the clock must meet its stability requirements. We plan to incorporate in-situ monitoring of the

final upconverted clocks on the AMC modules and will also periodically test the distributed signals at each crate to ensure that they have remained synchronized.

The  $\mu$ TCA crate will reside about one meter from the storage ring, where the residual magnetic field ranges from 30 to 60 Gauss. The crate, electronics, and power supply can potentially perturb the precision field, both statically and dynamically. The main concern for the static perturbation is the presence of ferromagnetic materials, which must be limited to several hundred grams at the proposed location. VadaTech has previous experience in migrating other  $\mu$ TCA chassis from their standard steel-based configuration to an Aluminum chassis. They will provide us with a custom Aluminum chassis for the full order and sent us preliminary versions of the chassis for magnetics characterization before filling the full order. They worked with us to identify and control other areas of the crate and modules that contain ferromagnetic materials. Two modifications will be made to the crate to reduce the magnetic field perturbations.

With the final  $\mu$ TCA crate version in hand, we will assess the magnetic field perturbations from the magnetization of the backend electronics using the fringe fields of the precision magnet that will be established at Fermilab. We will then use our OPERA 3D simulations to extrapolate observed perturbations in the test magnet to the behavior in the fringe field of the storage ring. We will also use pickup coils to assess the dynamical perturbations, and shield the crate as necessary to minimize these. If these mitigations are insufficient, we can always locate the crates somewhat farther from the storage ring, at the expense of modest increases in the signal cable length.

Longer term drifts in the frequency will in principle cancel in the  $\omega_a/\omega_p$  ratio since both the  $\omega_a$  and the magnetic field measurements are tied to the same master clock. Reliance upon this cancellation would, however, require some care in the procedure that weights the magnetic field measurements with the muon fill statistics. The GPS stabilization of the master clock will minimize this drift and therefore also minimize our need to rely upon strict cancellation of any time dependence in the ratio.

In the CDR, a significant cost risk (\$500,000) was associated with a higher digitization rate. Since then, the ADS5401 family of ADCs became available, affording an increase to an 800 MSPS digitization rate. With TI donating the ADCs to Cornell U., the risk decreases significantly and is associated only with the final production-level cost of the higher speed grade FPGAs. The tests performed at Cornell U. with an 800 MSPS digitization rate have shown to be successful, and there is no further cost risk associated with a higher digitization rate.

### 18.1.8 Quality Assurance

Cornell has established a  $\mu$ TCA test station to assess the performance of the  $\mu$ TCA platform, the AMC13 modules, and the WFD AMCs themselves. We are entering our third and last stage of prototyping. The initial stage gave us a one-channel design to verify the fundamental design and launch the majority of the firmware development, without facing the board density issue simultaneously. In the second stage, we moved to the full five-channel design with the denser component layout as well as migrating from the ADS5463 (500 MSPS) to the ADS5401 (800 MSPS) ADC. This prototype has undergone significant testing to assure that the baseline requirements for  $g-2$  are met. In the current third stage, we keep the

baseline five-channel design from the Revision 0.

We will produce enough of the second five-channel prototype to allow us to fully populate the  $\mu$ TCA crate, as planned for the experiment, so that we can ensure the entire system under full load can meet the specifications and that we do not encounter unanticipated cross talk or clock biasing with the full system.

Several months of burn-in and testing of the production WFD AMC modules will occur after their receipt. After an initial shakedown period to identify infant mortality, each module will be evaluated with an acceptance test suite that incorporates per-module and per-channel criteria. Some of the critical tests will include

- measurement of the phase stability and frequency upconversion factor at the part per billion level using the WFD AMC clock monitoring output,
- cross check of each channels upconversion factor to better than 10 parts per billion using a precision input sine wave,
- measurement of the linearity of each channel, and
- measurement of the noise level of each channel.

All results will be logged in a travelers database referenced by the WFD AMC serial number, which will form the basis of the modules MAC address. With production of the full complement of modules slated for fall 2015, we will have had ample time for testing and shakedown of the modules before installation at Fermilab the following summer.

The U. of Illinois group will test and evaluate each component of the clock system prior to installation. In addition, we will perform detailed in-situ timing for each path in the final experimental configuration. This “timing-in” is necessary to insure that the synchronization signals are delivered to each location simultaneously.

# References

- [1] R. Carey and J. Miller, E821 Muon  $g - 2$  Internal Note No. 165, 1 (1993).
- [2] G.W. Bennett *et al.*, Phys. Rev. D 73, 072003 (2006).
- [3] V. Tishchenko *et al.* [MuLan Collaboration], Phys. Rev. D 87, 052003 (2013) arXiv:1211.0960 [hep-ex].
- [4] J. Mans *et al.* [CMS Collaboration], *CMS Technical Design Report for the Phase 1 Upgrade of the Hadron Calorimeter*, CERN-LHCC-2012-015, CMS-TDR-10, <http://cds.cern.ch/record/1481837>.
- [5] M. Pesaresi *et al.*, *FC7 board*, <https://indico.cern.ch/event/300897/contribution/4/material/slides/1.pdf>.
- [6] *fmc-dio-5chttla FMC 5-channel Digital I/O module*, <http://www.ohwr.org/projects/fmc-dio-5chttla/wiki>.
- [7] *AMC13 development project*, <http://www.amc13.info>.
- [8] *Timing, Trigger and Control Systems for the LHC*, <http://ttc.web.cern.ch/TTC/intro.html>.
- [9] *AMC13 event builder*, [http://ohm.bu.edu/~hazen/CMS/AMC13/UpdatedDAQPath\\_2014-05-01.pdf](http://ohm.bu.edu/~hazen/CMS/AMC13/UpdatedDAQPath_2014-05-01.pdf).
- [10] *The IPbus Protocol*, [http://ohm.bu.edu/~chill90/ipbus/ipbus\\_protocol\\_v2.0.pdf](http://ohm.bu.edu/~chill90/ipbus/ipbus_protocol_v2.0.pdf).
- [11] N.T. Rider, J.P. Alexander, J.A. Dobbins, M.G. Billing, R.E. Meller, M.A. Palmer, D.P. Peterson, and C.R. Strohman, *Development of an X-Ray Beam Size Monitor with Single Pass Measurement Capability for CESR-TA*, Proceedings of the 2011 Particle Accelerator Conference (PAC 11), 687 (2011).